Beyond Countable Alphabets: An Extension of the Information-Spectrum Approach

Shengtian Yang*†, Thomas Honold‡†, and Zhaoyang Zhang‡†
*School of Information and Electronic Engineering, Zhejiang Gongshang University, Hangzhou 310018, China
†Zhejiang Provincial Key Laboratory of Information Processing, Communication and Networking, Hangzhou 310027, China
‡College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310027, China
Email: yangst@codlab.net, honold@zju.edu.cn, ning_ming@zju.edu.cn

arXiv:1602.00376v2 [cs.IT] 9 Jul 2016

***Abstract*—A general approach is established for deriving one-shot performance bounds for information-theoretic problems on general alphabets beyond countable alphabets. It is mainly based on the quantization idea and a novel form of "likelihood ratio". As an example, one-shot lower and upper bounds for random number generation from correlated sources on general alphabets are derived.**

I. Introduction

The information-spectrum approach, since its introduction by Han and Verdú [1], has become one of the most important tools of information theory. Similar to Shannon's classical approach (see e.g., [2]), the information-spectrum approach is mainly confined to countable alphabets, though in some cases, its extension to continuous alphabets is direct. The difficulty comes from two aspects:

1) It is difficult to formulate information measures and performance bounds for arbitrary alphabets in a unified way. The Radon-Nikodym derivative seems a good candidate, or a good basis for defining information measures, but it exists only if the absolute-continuity condition is satisfied.

2) It is difficult to extend certain useful proof techniques from finite alphabets to general alphabets, and even to countably infinite alphabets. For example, it is not easy to construct an analog of the random-bin map when the domain is uncountable, and it is also difficult to generalize structured random coding techniques (such as random matrices) to the case of infinite alphabets.

This paper will partly solve this problem by providing a general approach for one-shot performance bounds, an important part of the information-spectrum approach. Our main approach is based on quantization, which effectively overcomes the second difficulty. Given any problem on infinite alphabets, we take the following steps: 1) modifying the problem by quantization so that the modified version can be described on finite alphabets and hence has already a solution; 2) converting the solution of the modified problem into a solution of the original problem; and 3) repeating the first and second steps with a sequence of quantizations with increasing resolution to get the asymptotically optimal solution. This idea may look simple, but is difficult to be developed into a general approach. In the rest of this paper, we will give the basic results of this approach and then illustrate this approach by an example. We will also introduce a novel form of "likelihood ratio", a generalization of the Radon-Nikodym derivative that perfectly solves the first difficulty.

We close this section with some notations used throughout this paper. The field of real numbers is denoted by $\mathbf{R}$, and the set of integers from 1 to n is denoted by $[n]$. When performing probabilistic analysis, all objects of study are related to a basic probability space $(\Omega, \mathfrak{A}, P)$ with $\mathfrak{A}$ a σ-algebra in Ω and P a probability measure on $(\Omega, \mathfrak{A})$. A random element in a measurable space $(\mathcal{X}, \mathfrak{X})$ is a measurable mapping from Ω into $\mathcal{X}$. For probability measures μ and ν on $(\mathcal{X}, \mathfrak{X})$, the *statistical distance* between μ and ν is

$$\mathrm{d}(\mu, \nu) := \sup_{C \in \mathfrak{X}} |\mu(C) - \nu(C)|.$$

Measure-theoretic methods will be used in this paper. Readers not familiar with measure theory are referred to [3]–[5]. All measures considered in this paper are finite. For a topological space S, its Borel σ-algebra generated by the topology in S is denoted by $\mathfrak{B}(S)$. The integral of a real-valued, measurable function f on some measure space $(\mathcal{X}, \mathfrak{X}, \lambda)$ is denoted by $\lambda f = \lambda(f)$. When λ is a probability measure, we write $\mathrm{E}_\lambda f = \mathrm{E}_\lambda(f)$ in place of λf and write $\mathrm{E} f$ if $\lambda = P$. If $\lambda|f| < +\infty$, then λf induces a finite signed measure $(f \cdot \lambda)(A) := \lambda(f 1_A)$ on $(\mathcal{X}, \mathfrak{X})$. Given a sub σ-algebra $\mathfrak{F}$ of $\mathfrak{X}$, we denote by $\lambda^{\mathfrak{F}} f$ the density function $\mathrm{d}(f \cdot \lambda)|_{\mathfrak{F}} / \mathrm{d}\lambda|_{\mathfrak{F}}$, which is called the conditional expectation of f with respect to $\mathfrak{F}$ if λ is a probability measure. In this case, we wirte $\mathrm{E}_\lambda^{\mathfrak{F}} f$ in place of $\lambda^{\mathfrak{F}} f$. For $A \in \mathfrak{X}$, $\lambda^{\mathfrak{F}}(A) := \lambda^{\mathfrak{F}} 1_A$ defines a kernel from $(\mathcal{X}, \mathfrak{F})$ to $(\mathcal{X}, \mathfrak{X})$. For a kernel μ from $(\mathcal{X}, \mathfrak{X})$ to $(\mathcal{Y}, \mathfrak{Y})$, we denote by μ^B the real-valued map $x \mapsto \mu(x, B)$ for some fixed $B \in \mathfrak{Y}$, and by μ_x the measure on $(\mathcal{Y}, \mathfrak{Y})$ for some fixed $x \in \mathcal{X}$. Thus $\lambda(\mu^B)$, as a function of B, becomes a measure on $(\mathcal{Y}, \mathfrak{Y})$, and is usually written as $\lambda(\mu)$. The extended kernel $\overline{\mu}$ of μ is a kernel from $(\mathcal{X}, \mathfrak{X})$ to $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y})$ given by $(x, C) \mapsto \mu(x, C_x)$ with $C_x = \{y : (x, y) \in C\}$. The product $\lambda \times \mu := \lambda(\overline{\mu})$ is a measure on $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y})$, and we simply write $\lambda\mu$ when there is no possible ambiguity. Note that $\lambda\mu$ coincides with the product-measure notion when μ reduces to a measure on $(\mathcal{Y}, \mathfrak{Y})$.

II. The Quantization Approach

In this section, we will establish the basic results of the quantization approach. Because of the space limitation, most

simple proofs are omitted. The readers are referred to [6] for omitted proofs.

We first give an overview of the main tricks of our approach.

Trick 2.1 (Quantization): Given a measurable space $(\mathcal{X}, \mathfrak{X})$, a finite quantization of $\mathcal{X}$ can be characterized by a finite sub σ-algebra $\mathfrak{F}$ of $\mathfrak{X}$, which induces a natural projection $\pi_{\mathfrak{F}}$ from $\mathcal{X}$ to $\mathrm{atoms}(\mathfrak{F})$ given by $x \mapsto z$ such that $x \in z$, where $\mathrm{atoms}(\mathfrak{F})$ is the set of all elements in $\mathfrak{F}$ that cannot be decomposed into smaller pieces that are also in $\mathfrak{F}$. In fact, $\mathrm{atoms}(\mathfrak{F})$ forms a finite partition of $\mathcal{X}$.

Trick 2.2 (Approximation by Theorems 2.4, 2.6, 2.15, and Corollary 2.5): Let $f = (f_i)_{i=1}^{\ell}$ be a family of real-valued integrable functions on the measure space $(\mathcal{X}_1 \times \mathcal{X}_2, \mathfrak{X}_1 \times \mathfrak{X}_2, \mu)$. Let $\epsilon > 0$. By Theorem 2.4, for each $i \in [\ell]$, there is a finite sub σ-algebra $\mathfrak{C}_{i,j}$ of $\mathfrak{X}_j$ for each $j \in [2]$ such that

$$\mu|\mu^{\mathfrak{F}} f_i - f_i| < \epsilon$$

for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C}_{i,1} \times \mathfrak{C}_{i,2} \subseteq \mathfrak{F} \subseteq \mathfrak{X}_1 \times \mathfrak{X}_2$. Taking

$$\mathfrak{D}_j = \sigma\left(\bigcup_{i\in[\ell]} \mathfrak{C}_{i,j}\right)$$

and $\mathfrak{F} = \mathfrak{D}_1 \times \mathfrak{D}_2$, we thus have

$$\mu|\mu^{\mathfrak{F}} f_i - f_i| < \epsilon$$

for all $i \in [\ell]$. In the same vein and by Corollary 2.5, we can show that, for any $\nu \ll \mu$, there is a finite sub σ-algebra $\mathfrak{F}$ of $\mathfrak{X}_1 \times \mathfrak{X}_2$ such that

$$\nu\{|\mu^{\mathfrak{F}} f_i - f_i| \geq \epsilon\} < \epsilon$$

for all $i \in [\ell]$. Then we can find a sequence $(\mathfrak{F}_n)_{n=1}^{\infty}$ of finite sub σ-algebras so that each sequence $g_i^{(n)} = \mu^{\mathfrak{F}_n} f_i$ converges in ν-measure to f_i, or further,

$$\|(g_i^{(n)})_{i=1}^{\ell} - (f_i)_{i=1}^{\ell}\|_p \xrightarrow{\nu} 0$$

for any p-norm of $\mathbf{R}^{\ell}$ with $p \geq 1$, where $\xrightarrow{\nu}$ is the shorthand of convergence in ν-measure. Similar tricks with Theorems 2.6 and 2.15 also work for statistical distances and likelihood ratios (Trick 2.3).

Trick 2.3 (Handling likelihood ratios by Theorems 2.15): Likelihood ratios may be the objects most often occurring in a one-shot performance bound. Let μ and ν be two measures on $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y})$. The likelihood ratio of μ to ν is

$$[\mathrm{d}\mu : \mathrm{d}\nu] = [\mathrm{d}\mu/\mathrm{d}(\mu+\nu) : \mathrm{d}\nu/\mathrm{d}(\mu+\nu)] \quad \text{(Definition 2.14)},$$

a $(\mathbb{P}, \mathfrak{B}(\mathbb{P}))$-valued measurable function on $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y})$, where $\mathbb{P}$ is the half projective line defined by Definitions 2.9 and 2.10. For any $\xi \ll \mu+\nu$, we can find a sequence $(\mathfrak{F}_n)_{n=1}^{\infty}$ of finite sub σ-algebras of $\mathfrak{X} \times \mathfrak{Y}$ such that

$$d_{\mathbb{P}}([\mathrm{d}\mu|_{\mathfrak{F}_n} : \mathrm{d}\nu|_{\mathfrak{F}_n}], [\mathrm{d}\mu : \mathrm{d}\nu]) \xrightarrow{\xi} 0$$
(Theorem 2.15 and Trick 2.2).

Having introduced the main tricks, we proceed to introduce the details of the quantization approach.

Theorem 2.4: Let f be a real-valued integrable function on the measure space $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y}, \mu)$. Then for any $\epsilon > 0$, there is a finite sub σ-algebra $\mathfrak{C}$ of $\mathfrak{X}$ and a finite sub σ-algebra $\mathfrak{D}$ of $\mathfrak{Y}$ such that

$$\mu|\mu^{\mathfrak{F}} f - f| < \epsilon \tag{1}$$

for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$.

Proof: We say that a real-valued integrable function f on $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y}, \mu)$ can be finitely approximated if for any $\epsilon > 0$, there are a finite sub σ-algebra $\mathfrak{C}$ of $\mathfrak{X}$ and a finite sub σ-algebra $\mathfrak{D}$ of $\mathfrak{Y}$ such that $\mu|\mu^{\mathfrak{F}} f - f| < \epsilon$ for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$. We define

$$\mathcal{H} = \{f : f \text{ is integrable and can be finitely approximated}\}$$

and $\mathcal{A} = \{C \times D : C \in \mathfrak{X}, D \in \mathfrak{Y}\}$. It is clear that $\mathcal{A}$ is a π-system containing $\mathcal{X} \times \mathcal{Y}$, and we have:

(a) If $A = C \times D \in \mathcal{A}$, then $1_A \in \mathcal{H}$ with $\mathfrak{C} = \{\varnothing, C, C^{\mathsf{c}}, \mathcal{X}\}$ and $\mathfrak{D} = \{\varnothing, D, D^{\mathsf{c}}, \mathcal{Y}\}$.

(b) If $f, g \in \mathcal{H}$, then there are finite sub σ-algebras $\mathfrak{C}'$, $\mathfrak{C}''$ of $\mathfrak{X}$ and $\mathfrak{D}'$, $\mathfrak{D}''$ of $\mathfrak{Y}$ such that $\mu|\mu^{\mathfrak{F}'} f - f| < \epsilon/2$ and $\mu|\mu^{\mathfrak{F}''} g - g| < \epsilon/2$ for every σ-algebra $\mathfrak{C}' \times \mathfrak{D}' \subseteq \mathfrak{F}' \subseteq \mathfrak{X} \times \mathfrak{Y}$ and every σ-algebra $\mathfrak{C}'' \times \mathfrak{D}'' \subseteq \mathfrak{F}'' \subseteq \mathfrak{X} \times \mathfrak{Y}$, respectively. Then

$$\mu|\mu^{\mathfrak{F}}(f+g) - (f+g)| \leq \mu|\mu^{\mathfrak{F}} f - f| + \mu|\mu^{\mathfrak{F}} g - g| < \epsilon$$

for every σ-algebra $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$ with $\mathfrak{C} = \sigma(\mathfrak{C}' \cup \mathfrak{C}'')$ and $\mathfrak{D} = \sigma(\mathfrak{D}' \cup \mathfrak{D}'')$. In other words, $f + g \in \mathcal{H}$. In a similar way, we can show that $cf \in \mathcal{H}$ for $c \in \mathbf{R}$.

(c) If $f_n \in \mathcal{H}$ converges everywhere to an integrable function g (including all bounded functions for finite μ) with $|f_n| \leq |g|$, then by the dominated convergence theorem, $\mu|f_N - g| < \epsilon/4$ for some large integer N. Furthermore, since $f_N \in \mathcal{H}$, there are finite σ-algebras $\mathfrak{C}$ of $\mathfrak{X}$ and $\mathfrak{D}$ of $\mathfrak{Y}$ such that $\mu|\mu^{\mathfrak{F}} f_N - f_N| < \epsilon/2$ for every σ-algebra $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$, so that

$$\begin{aligned}\mu|\mu^{\mathfrak{F}} g - g| &\leq \mu|\mu^{\mathfrak{F}} g - \mu^{\mathfrak{F}} f_N| + \mu|\mu^{\mathfrak{F}} f_N - f_N| + \mu|f_N - g| \\ &\leq \mu|g - f_N| + \mu|\mu^{\mathfrak{F}} f_N - f_N| + \mu|f_N - g| \\ &= 2\mu|f_N - g| + \mu|\mu^{\mathfrak{F}} f_N - f_N| < \epsilon\end{aligned}$$

for every σ-algebra $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$. Therefore, $g \in \mathcal{H}$.

By the monotone class theorem for functions ([4, Theorem 6.1.3]) with properties (a)–(c), we conclude that $\mathcal{H}$ contains all bounded functions measurable with respect to $\sigma(\mathcal{A}) = \mathfrak{X} \times \mathfrak{Y}$. Again by (c) with $f_n(p) = g(p)1\{g(p) \leq n\}$ for arbitrary integrable g, it is easy to see that $\mathcal{H}$ contains all integrable functions. □

Corollary 2.5: Let f be a real-valued integrable function on the measure space $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y}, \mu)$. Let ν be a measure such that $\nu \ll \mu$. Then for any $\epsilon > 0$, there is a finite sub σ-algebra $\mathfrak{C}$ of $\mathfrak{X}$ and a finite sub σ-algebra $\mathfrak{D}$ of $\mathfrak{Y}$ such that

$$\nu\{|\mu^{\mathfrak{F}} f - f| \geq \epsilon\} < \epsilon$$

for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$.

Theorem 2.6: Let μ be a probability kernel from $(\mathcal{X}, \mathfrak{X}, \lambda)$ to $(\mathcal{Y}, \mathcal{P}(\mathcal{Y}))$ with $\mathcal{Y}$ at most countable. Then for any $\epsilon > 0$, there is a finite sub σ-algebra $\mathfrak{C}$ of $\mathfrak{X}$ such that

$$\mathrm{E}_\lambda \, \mathrm{d}(\mathrm{E}_\lambda^{\mathfrak{F}} \mu, \mu) < \epsilon$$

for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C} \subseteq \mathfrak{F} \subseteq \mathfrak{X}$.

A performance bound obtained by the information-spectrum approach is often expressed in terms of some kind of random likelihood ratios, it is thus necessary to understand the notions of likelihood ratios well. Usually, a likelihood ratio is expressed as a Radon-Nikodym derivative of two probability measures. Sometimes, however, we will encounter a more complicated form of likelihood ratios, a Radon-Nikodym derivative of two probability kernels. In the discrete case, it can be written as $P_{Y|X}(y \mid x)/P_{\hat{Y}|X}(y \mid x)$, and we have

$$\frac{P_{Y|X}(y \mid x)}{P_{\hat{Y}|X}(y \mid x)} = \frac{P_{XY}(x, y)}{P_{X\hat{Y}}(x, y)}.$$

A generalization of this identity is given as follows.

Theorem 2.7: Let μ and ν be two kernels from $(\mathcal{X}, \mathfrak{X}, \lambda)$ to $(\mathcal{Y}, \mathfrak{Y})$ such that $\mu_x \ll \nu_x$ for λ-almost every x in $\mathcal{X}$. Then $\lambda\mu \ll \lambda\nu$ and for λ-almost every x,

$$\frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}(x, y) = \frac{\mathrm{d}\mu_x}{\mathrm{d}\nu_x}(y)$$

ν_x-almost everywhere. If $\mathrm{d}\mu_x/\mathrm{d}\nu_x$ has a version, say $f(x, y)$, that is measurable with respect to $\mathfrak{X} \times \mathfrak{Y}$, then $f = \mathrm{d}(\lambda\mu)/\mathrm{d}(\lambda\nu)$ $\lambda\nu$-almost everywhere.

This theorem tells us that in general cases we need to use the form $\mathrm{d}(\lambda\mu)/\mathrm{d}(\lambda\nu)$ in place of $\mathrm{d}\mu_x/\mathrm{d}\nu_x$ because the former is always measurable with respect to $\mathfrak{X} \times \mathfrak{Y}$. A useful consequence of Theorem 2.7 is:

Corollary 2.8: Let μ and ν be two probability kernels from the probability space $(\mathcal{X}, \mathfrak{X}, \lambda)$ to $(\mathcal{Y}, \mathfrak{Y})$. Then $\mathrm{d}(\lambda\mu, \lambda\nu) = \mathrm{E}_\lambda \, \mathrm{d}(\mu_x, \nu_x)$.

The Radon-Nikodym derivative $\mathrm{d}\mu/\mathrm{d}\nu$ cannot handle all cases of likelihood ratios, because it does not exist if μ is not absolutely continuous with respect to ν. In this case, we need a more general form of likelihood ratios based on the approach of [7].

Recall that the (real) projective line is defined as the set of lines through the origin in the affine plane $\mathbf{R}^2$, and points $\mathbf{R}(x, y)$ of the projective line are written as $(x : y)$ (homogeneous coordinates), reflecting the fact that $\mathbf{R}(x, y) = \mathbf{R}(z, w)$ iff $x/y = z/w$ (for $y, w \neq 0$). In analogy, we define $\mathbb{P}$, the *nonnegative part* of the projective line, as follows:

Definition 2.9: A pair $(r, s) \in \mathbf{R}^2$ is said to be admissible if $(r, s) \in \mathbf{R}^2_{\geq 0} \setminus \{(0, 0)\}$. The set $\mathbf{R}_{>0}(r, s) := \{t(r, s) : t > 0\}$ forms a ray through the origin in $\mathbf{R}^2_{\geq 0}$ iff (r, s) is admissible. The *half projective line* $\mathbb{P}$ is defined as the set of rays through the origin in $\mathbf{R}^2_{\geq 0}$, and points $\mathbf{R}_{>0}(r, s)$ of $\mathbb{P}$ are written as $[r : s]$, or simply as r/s when $s \neq 0$ and there is no possible ambiguity. The natural projection $\pi_{\mathbb{P}} : \mathbf{R}^2_{\geq 0} \setminus \{(0, 0)\} \to \mathbb{P}$ given by $(x, y) \mapsto [x : y]$ thus induces a quotient topology in $\mathbb{P}$, so that $\pi_{\mathbb{P}}$ becomes a quotient map and is measurable with respect to the corresponding Borel σ-algebras.

Since the map $\rho : \mathbf{R}^2_{\geq 0} \setminus \{(0, 0)\} \to [0, 1]$ given by $(x, y) \mapsto x/(x+y)$ is a quotient map, it follows from [8, Corollary 22.3] that ρ induces a homeomorphism $\kappa : \mathbb{P} \to [0, 1]$ given by

$$[x : y] \mapsto \frac{x}{x + y},$$

which further induces a metric and an order on $\mathbb{P}$.

Definition 2.10: The metric $d_{\mathbb{P}}$ on $\mathbb{P}$ is defined by

$$\begin{aligned} d_{\mathbb{P}}([r_1 : s_1], [r_2 : s_2]) &:= |\kappa([r_1 : s_1]) - \kappa([r_2 : s_2])| \\ &= \left| \frac{r_1}{r_1 + s_1} - \frac{r_2}{r_2 + s_2} \right|. \end{aligned}$$

The order $\leq$ of $\mathbb{P}$ is defined by

$$\begin{aligned} [r_1 : s_1] \leq [r_2 : s_2] &\Leftrightarrow \kappa([r_1 : s_1]) \leq \kappa([r_2 : s_2]) \\ &\Leftrightarrow r_1 s_2 - r_2 s_1 \leq 0. \end{aligned}$$

It is clear that $\mathbb{P}$ with metric $d_{\mathbb{P}}$ is a complete separable metric space.

For any real-valued functions f and g on $\mathcal{X}$, if $(f(x), g(x))$ is admissible for all $x \in \mathcal{X}$, then the function $[f : g](x) := [f(x) : g(x)]$ is well defined and is also called admissible (on $\mathcal{X}$). Conversely, any $\mathbb{P}$-valued function on $\mathcal{X}$ can be written as $[f : g]$ with f and g two real-valued functions on $\mathcal{X}$. Below are some properties of $\mathbb{P}$-valued functions.

Proposition 2.11: The $(\mathbb{P}, \mathfrak{B}(\mathbb{P}))$-valued function $[f : g]$ on $(\mathcal{X}, \mathfrak{X})$ is measurable if f and g are both measurable.

Proposition 2.12: If $[f_1 : g_1] = [f_2 : g_2]$ with f_1, f_2, g_1, and g_2 all real-valued measurable functions on $\mathcal{X}$, then there is a real-valued measurable function t on $\mathcal{X}$ such that $t(x) \neq 0$, $f_1(x) = t(x) f_2(x)$, and $g_1(x) = t(x) g_2(x)$ for all $x \in \mathcal{X}$.

Let μ be a measure on $(\mathcal{X}, \mathfrak{X})$. If $[f : g]$ is admissible on $\mathcal{X}$ except a μ-negligible set of points, then we say $[f : g]$ is admissible μ-almost everywhere. Similarly, if $[f_1 : g_1] = [f_2 : g_2]$ is true for all $x \in \mathcal{X}$ except a μ-negligible set of points, we say $[f_1 : g_1] = [f_2 : g_2]$ μ-almost everywhere.

Proposition 2.13: Let f and g be two real-valued integrable functions on the measure space $(\mathcal{X}, \mathfrak{X}, \mu)$. If $[f : g]$ is admissible μ-almost everywhere, then the conditional expectation $\mu^{\mathfrak{F}}[f : g] := [\mu^{\mathfrak{F}} f : \mu^{\mathfrak{F}} g]$ with respect to some sub σ-algebra $\mathfrak{F}$ is also admissible μ-almost everywhere.

We are now ready to define the general form of likelihood ratios.

Definition 2.14: Let μ and ν be two measures on $(\mathcal{X}, \mathfrak{X})$. The *likelihood ratio* $[\mathrm{d}\mu : \mathrm{d}\nu]$ of μ to ν is defined to be $[\mathrm{d}\mu/\mathrm{d}(\mu + \nu) : \mathrm{d}\nu/\mathrm{d}(\mu + \nu)]$, which is admissible $(\mu + \nu)$-almost everywhere.

Likelihood ratios enjoy the following property, which is an easy consequence of Corollary 2.5.

Theorem 2.15: Let μ and ν be two measures on $(\mathcal{X} \times \mathcal{Y}, \mathfrak{X} \times \mathfrak{Y})$. Let ξ be a measure such that $\xi \ll \mu + \nu$. Then for any $\epsilon > 0$, there is a finite sub σ-algebra $\mathfrak{C}$ of $\mathfrak{X}$ and a finite sub σ-algebra $\mathfrak{D}$ of $\mathfrak{Y}$ such that

$$\xi\{d_{\mathbb{P}}([\mathrm{d}\mu|_{\mathfrak{F}} : \mathrm{d}\nu|_{\mathfrak{F}}], [\mathrm{d}\mu : \mathrm{d}\nu]) \geq \epsilon\} < \epsilon$$

for every σ-algebra satisfying $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$.

III. An Example: Separate Random Number Generation from Correlated Sources

In this section, we will explain the quantization approach by an example: separate random number generation from correlated sources. For its importance in information theory, the readers are referred to [9]. The finite-alphabet case of this problem has been extensively studied in [10] and the references therein. We will now extend this result to the case of general alphabets.

We first briefly introduce the problem of separate random number generation. For simplicity, we only consider the case of two correlated sources with side information at the tester.

Let $X = (X_0, X_1, X_2)$ be a triple of correlated random elements in $\mathcal{X}_0 \times \mathcal{X}_1 \times \mathcal{X}_2$. Let $\varphi = (\varphi_1, \varphi_2)$ be a pair of (randomness) extractors $\mathcal{X}_i \to \mathcal{Y}_i$ with $\mathcal{Y}_i$ finite ($i = 1, 2$). We are interested in the minimum value of the statistical distance

$$d(X \mid \varphi) := \mathrm{d}(P_{X_0\varphi_1(X_1)\varphi_2(X_2)}, P_{X_0}\mathrm{U}_{\mathcal{Y}_1}\mathrm{U}_{\mathcal{Y}_2})$$

over all pairs φ of extractors, where $\mathrm{U}_{\mathcal{Y}_1}$ and $\mathrm{U}_{\mathcal{Y}_2}$ denote the uniform distributions over $\mathcal{Y}_1$ and $\mathcal{Y}_2$, respectively. The next theorem gives one-shot bounds of $d(X \mid \varphi)$ in the case of finite alphabets.

Theorem 3.1 ([10]): Let $X = (X_0, X_1, X_2)$ be a triple of correlated random elements in a finite product alphabet $\mathcal{X}_0 \times \mathcal{X}_1 \times \mathcal{X}_2$.

1) For $r > 1$, there exists a pair φ of extractors such that

$$d(X \mid \varphi) \le P\{(T^i_X(X))^3_{i=1} \notin A_r\} + \frac{\sqrt{3}}{2} r^{-1/2},$$

where

$$\begin{aligned} T^1_X(x_0, x_1, x_2) &:= \frac{1}{P_{X_1|X_0}(x_1 \mid x_0)}, \\ T^2_X(x_0, x_1, x_2) &:= \frac{1}{P_{X_2|X_0}(x_2 \mid x_0)}, \\ T^3_X(x_0, x_1, x_2) &:= \frac{1}{P_{X_1X_2|X_0}(x_1, x_2 \mid x_0)}, \end{aligned}$$

$A_r := I_{r|\mathcal{Y}_1|} \times I_{r|\mathcal{Y}_2|} \times I_{r|\mathcal{Y}_1||\mathcal{Y}_2|}$, and $I_t := (t, +\infty)$.

2) Conversely, every pair φ of extractors satisfies

$$d(X \mid \varphi) \ge P\{(T^i_X(X))^3_{i=1} \notin A_r\} - 3r$$

for all $0 < r < 1$.

Now let us prove a general-alphabet version of Theorem 3.1.

Theorem 3.2: Let $X = (X_0, X_1, X_2)$ be a triple of correlated random elements in $(\mathcal{X}_0 \times \mathcal{X}_1 \times \mathcal{X}_2, \mathfrak{X}_0 \times \mathfrak{X}_1 \times \mathfrak{X}_2)$.[1]

1) For $r > 1$ and $\epsilon > 0$, there exists a pair φ of extractors such that

$$d(X \mid \varphi) \le P\{(T^i_X(X))^3_{i=1} \notin A_r\} + \frac{\sqrt{3}}{2} r^{-1/2} + \epsilon, \quad (2)$$

where

$$\begin{aligned} T^1_X(x_0, x_1, x_2) &:= [\mathrm{d}P_{X_0X_1X_1} : \mathrm{d}P_{X_0X_1}P_{X_1|X_0}](x_0, x_1, x_1), \\ T^2_X(x_0, x_1, x_2) &:= [\mathrm{d}P_{X_0X_2X_2} : \mathrm{d}P_{X_0X_2}P_{X_2|X_0}](x_0, x_2, x_2), \\ T^3_X(x_0, x_1, x_2) &:= [\mathrm{d}P_{X_0X_1X_2X_1X_2} \\ &\qquad : \mathrm{d}P_{X_0X_1X_2}P_{X_1X_2|X_0}](x_0, x_1, x_2, x_1, x_2), \end{aligned}$$

$A_r := I_{r|\mathcal{Y}_1|} \times I_{r|\mathcal{Y}_2|} \times I_{r|\mathcal{Y}_1||\mathcal{Y}_2|}$, and $I_t := ([t : 1], [1 : 0])$.

2) Conversely, every pair φ of extractors satisfies

$$d(X \mid \varphi) \ge P\{(T^i_X(X))^3_{i=1} \notin A_r\} - 3r$$

for all $0 < r < 1$.

Proof: When the alphabets are all finite, it is clear that

$$\begin{aligned} T^1_X(x_0, x_1, x_2) &= [1 : P_{X_1|X_0}(x_1 \mid x_0)] \\ T^2_X(x_0, x_1, x_2) &= [1 : P_{X_2|X_0}(x_2 \mid x_0)] \\ T^3_X(x_0, x_1, x_2) &= [1 : P_{X_1X_2|X_0}(x_1, x_2 \mid x_0)] \end{aligned}$$

P_X-almost everywhere, and thus the theorem is trivially true because of Theorem 3.1.

1) Direct part: We first show that the direct part is true for general $(\mathcal{X}_0, \mathfrak{X}_0)$ and finite $\mathcal{X}_i$ for $i \in [2]$. By Trick 2.2 with Theorems 2.6 and 2.15, we can find a sequence $(\mathfrak{F}_n)_{n=1}^\infty$ of finite sub σ-algebras of $\mathfrak{X}_0$ such that

$$\lim_{n\to\infty} \mathrm{E}_{P_{X_0}} \mathrm{d}(P_{X_1X_2|X_0=x_0}, P_{X_1X_2|Z_{n,0}=\pi_{\mathfrak{F}_n}(x_0)}) = 0 \quad (3)$$

and

$$d_{\mathbb{P}}(T^i_X, T^i_{g_n(X)} \circ g_n) \xrightarrow{P_X} 0 \quad (4)$$

for $i \in [3]$, where $g_n(x_0, x_1, x_2) = (\pi_{\mathfrak{F}_n}(x_0), x_1, x_2)$ and $Z_{n,0} = \pi_{\mathfrak{F}_n}(X_0)$. By Theorem 3.1, there is a pair $\varphi_n = (\varphi_{n,1}, \varphi_{n,2})$ of extractors such that

$$d(g_n(X) \mid \varphi_n) \le P\{(T^i_{g_n(X)}(g_n(X)))^3_{i=1} \notin A_r\} + \frac{\sqrt{3}}{2} r^{-1/2}$$

and it follows from Corollary 2.8 that

$$\begin{aligned} &\big|d(X \mid \varphi_n) - d(g_n(X) \mid \varphi_n)\big| \\ &= \big|\mathrm{E}_{P_{X_0}} \mathrm{d}\left(P_{\varphi_{n,1}(X_1)\varphi_{n,2}(X_2)|X_0=x_0}, \mathrm{U}_{\mathcal{Y}_1}\mathrm{U}_{\mathcal{Y}_2}\right) \\ &\quad - \mathrm{E}_{P_{X_0}} \mathrm{d}\left(P_{\varphi_{n,1}(X_1)\varphi_{n,2}(X_2)|Z_{n,0}=\pi_{\mathfrak{F}_n}(x_0)}, \mathrm{U}_{\mathcal{Y}_1}\mathrm{U}_{\mathcal{Y}_2}\right)\big| \\ &\le \mathrm{E}_{P_{X_0}} \big|\mathrm{d}\left(P_{\varphi_{n,1}(X_1)\varphi_{n,2}(X_2)|X_0=x_0}, \mathrm{U}_{\mathcal{Y}_1}\mathrm{U}_{\mathcal{Y}_2}\right) \\ &\quad - \mathrm{d}\left(P_{\varphi_{n,1}(X_1)\varphi_{n,2}(X_2)|Z_{n,0}=\pi_{\mathfrak{F}_n}(x_0)}, \mathrm{U}_{\mathcal{Y}_1}\mathrm{U}_{\mathcal{Y}_2}\right)\big| \\ &\le \mathrm{E}_{P_{X_0}} \mathrm{d}\big(P_{\varphi_{n,1}(X_1)\varphi_{n,2}(X_2)|X_0=x_0}, \\ &\quad P_{\varphi_{n,1}(X_1)\varphi_{n,2}(X_2)|Z_{n,0}=\pi_{\mathfrak{F}_n}(x_0)}\big) \\ &\le \mathrm{E}_{P_{X_0}} \mathrm{d}\big(P_{X_1X_2|X_0=x_0}, P_{X_1X_2|Z_{n,0}=\pi_{\mathfrak{F}_n}(x_0)}\big) = \mathrm{o}(1), \end{aligned}$$

so that

$$\limsup_{n\to\infty} d(X \mid \varphi_n) \le P\{(T^i_X(X))^3_{i=1} \notin A_r\} + \frac{\sqrt{3}}{2} r^{-1/2}$$

by the Portmanteau theorem [5, Theorem 3.25] with (4), and therefore φ_n satisfies (2) for sufficiently large n.

We are now ready to prove the general case. By Trick 2.2 with Theorem 2.15, we can find a sequence $((\mathfrak{C}_{n,1}, \mathfrak{C}_{n,2}))_{n=1}^\infty$ of pairs of finite sub σ-algebras such that, for all $i \in [3]$,

$$d_{\mathbb{P}}(T^i_X, T^i_{h_n(X)} \circ h_n) \xrightarrow{P_X} 0, \quad (5)$$

where $h_n(x_0, x_1, x_2) = (x_0, \pi_{\mathfrak{C}_{n,1}}(x_1), \pi_{\mathfrak{C}_{n,2}}(x_2))$. Then, for every n, there is a pair $\psi_n = (\psi_{n,1}, \psi_{n,2})$ of extractors such that

$$\begin{aligned} d(h_n(X) \mid \psi_n) \le{}& P\{(T^i_{h_n(X)}(h_n(X)))^3_{i=1} \notin A_r\} \\ &+ \frac{\sqrt{3}}{2} r^{-1/2} + \frac{\epsilon}{2}. \end{aligned}$$

[1] It is assumed that the σ-algebra of a countable alphabet is its power set.

Let $\varphi_n = (\psi_{n,1} \circ \pi_{\mathfrak{C}_{n,1}}, \psi_{n,2} \circ \pi_{\mathfrak{C}_{n,2}})$. We further have

$$\begin{aligned}
&\limsup_{n\to\infty} d(X \mid \varphi_n) \\
&= \limsup_{n\to\infty} d(h_n(X) \mid \psi_n) \\
&\le \limsup_{n\to\infty} P\{(T^i_{h_n(X)}(h_n(X)))_{i=1}^3 \notin A_r\} + \frac{\sqrt{3}}{2} r^{-1/2} + \frac{\epsilon}{2} \\
&\le P\{(T^i_X(X))_{i=1}^3 \notin A_r\} + \frac{\sqrt{3}}{2} r^{-1/2} + \frac{\epsilon}{2},
\end{aligned}$$

where the last inequality follows from (5) and the Portmanteau theorem, and therefore φ_n satisfies (2) for sufficiently large n.

2) Converse part: Similar to the proof of direct part, we first prove the converse part with general $(\mathcal{X}_0, \mathfrak{X}_0)$ and other alphabets finite. Again by Trick 2.2 with Theorems 2.6 and 2.15, we have (3) and (4), so that

$$\begin{aligned}
d(X \mid \varphi) &= \lim_{n\to\infty} d(g_n(X) \mid \varphi) \\
&\ge \liminf_{n\to\infty} P\{(T^i_{g_n(X)}(g_n(X)))_{i=1}^3 \notin \overline{A_r}\} - 3r \\
&\ge P\{(T^i_X(X))_{i=1}^3 \notin \overline{A_r}\} - 3r,
\end{aligned}$$

where the last inequality follows from the Portmanteau theorem, and $\overline{A_r}$ denotes the closure of A_r.

Then we turn to the general case. By Trick 2.2 with Theorem 2.15, we can find a sequence $((\mathfrak{C}_{n,1}, \mathfrak{C}_{n,2}))_{n=1}^\infty$ of pairs of finite sub σ-algebras satisfying (5) and $\sigma(\varphi_i) \subseteq \mathfrak{C}_{n,i}$ for all n and i. Then for every n and i, we have $\varphi_i = \psi_{n,i} \circ \pi_{\mathfrak{C}_{n,i}}$ for some $\psi_{n,i}$, and therefore

$$\begin{aligned}
d(X \mid \varphi) &= \lim_{n\to\infty} d(h_n(X) \mid \psi_n) \\
&\ge \liminf_{n\to\infty} P\{(T^i_{g_n(X)}(g_n(X)))_{i=1}^3 \notin \overline{A_r}\} - 3r \\
&\ge P\{(T^i_X(X))_{i=1}^3 \notin \overline{A_r}\} - 3r,
\end{aligned}$$

where the last inequality follows from the Portmanteau theorem. Finally,

$$\begin{aligned}
d(X \mid \varphi) &\ge \lim_{k\to\infty} \left(P\{(T^i_X(X))_{i=1}^3 \notin \overline{A_{s_k}}\} - 3s_k\right) \\
&\ge P\left\{(T^i_X(X))_{i=1}^3 \notin \bigcup_{k=1}^\infty \overline{A_{s_k}}\right\} - 3r \\
&= P\{(T^i_X(X))_{i=1}^3 \notin A_r\} - 3r,
\end{aligned}$$

where $s_k = (r + 1/k) \wedge 1$, namely, the minimum of $r + 1/k$ and 1. □

Remark 3.3: Note that even if the alphabets $\mathcal{X}_i$ are all continuous, there are many nontrivial situations in which $T^i_X(X)$ does not degenerate to the point $[1:0]$. For simplicity, let us assume that X_2 is constant. If X_1 is a $(p, 1-p)$-mixture of a real number and a uniform random variable in $[0,1]$, both independent of X_0, then $T^3_X(X)$ is a random variable taking values in $\{1/p, [1:0]\}$ with probabilities p and $1-p$, respectively. Note that the extractors in our problem are fixed-length extractors. This example tells us that the performance of a fixed-length extractor is dominated by the worst case. If X_1 takes values in $\{X_0 - 0.5, X_0 + 0.5\}$ with equal probabilities and X_0 is uniformly distributed over $[0,1]$, then $T^3_X(X) = 2$ almost surely. This example shows that a random variable even with continuous distribution does not necessarily have infinite randomness because of the side information at the tester.

The one-shot bounds provided by Theorem 3.2 are tight enough for the first- and second-order asymptotic analysis. If we define

$$\ln[x:y] := \ln(x) - \ln(y) \in [-\infty, +\infty]$$

(which is well defined because $(x,y) \neq (0,0)$), then we can easily obtain a generalization of the achievable rate region in the finite case, with for example the *spectral inf-entropy rate* $\underline{\mathcal{H}}(X_1 \mid X_0)$ of general source X_1 given general source X_0 defined by

$$\text{p-}\liminf_{n\to\infty} \frac{1}{n} \ln\left[P_{X_0^{(n)} X_1^{(n)} X_1^{(n)}} : P_{X_0^{(n)} X_1^{(n)}} P_{X_1^{(n)}|X_0^{(n)}}\right].$$

IV. The Proofs of Results in Section II

Proof of Corollary 2.5: Since $\nu \ll \mu$, it follows from [3, Lemma 4.2.1] that there is a positive δ such that each measurable set A satisfying $\mu(A) < \delta$ also satisfies $\nu(A) < \epsilon$. Using Theorem 2.4 with $\epsilon' = \delta\epsilon/2$, we obtain by Markov's inequality that there is a finite sub σ-algebra $\mathfrak{C}$ of $\mathfrak{X}$ and a finite sub σ-algebra $\mathfrak{D}$ of $\mathfrak{Y}$ such that

$$\mu\{|\mu^{\mathfrak{F}} f - f| \ge \epsilon\} \le \frac{\mu|\mu^{\mathfrak{F}} f - f|}{\epsilon} = \frac{\delta}{2}$$

for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C} \times \mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X} \times \mathfrak{Y}$, so that $\nu\{|\mu^{\mathfrak{F}} f - f| \ge \epsilon\} < \epsilon$. □

Proof of Theorem 2.6: Since $\mathrm{E}_\lambda(\mu)$ is a probability measure on $\mathcal{Y}$, there is a finite subset B of $\mathcal{Y}$ such that

$$(\mathrm{E}_\lambda(\mu))(B^{\mathsf{c}}) = \mathrm{E}_\lambda(\mu^{B^{\mathsf{c}}}) < \frac{\epsilon}{2}.$$

For every $y \in B$, it follows from Theorem 2.4 that there is a finite sub σ-algebra $\mathfrak{C}_y$ of $\mathfrak{X}$ such that

$$\mathrm{E}_\lambda \left|\mathrm{E}^{\mathfrak{F}}_\lambda \mu^{\{y\}} - \mu^{\{y\}}\right| < \frac{\epsilon}{|B|}.$$

for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C}_y \subseteq \mathfrak{F} \subseteq \mathfrak{X}$. Let $\mathfrak{C} = \sigma(\bigcup_{y\in B} \mathfrak{C}_y)$, and then for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C} \subseteq \mathfrak{F} \subseteq \mathfrak{X}$,

$$\begin{aligned}
&\mathrm{E}_\lambda\, \mathrm{d}(\mathrm{E}^{\mathfrak{F}}_\lambda \mu, \mu) \\
&= \int \frac{1}{2} \sum_{y\in\mathcal{Y}} |\mathrm{E}^{\mathfrak{F}}_\lambda \mu(x, \{y\}) - \mu(x, \{y\})| \lambda(\mathrm{d}x) \\
&\le \int \frac{1}{2} \sum_{y\in B} |\mathrm{E}^{\mathfrak{F}}_\lambda \mu(x, \{y\}) - \mu(x, \{y\})| \lambda(\mathrm{d}x) \\
&\quad + \int \frac{1}{2} (\mathrm{E}^{\mathfrak{F}}_\lambda \mu(x, B^{\mathsf{c}}) + \mu(x, B^{\mathsf{c}})) \lambda(\mathrm{d}x) \\
&= \frac{1}{2} \sum_{y\in B} \mathrm{E}_\lambda \left|\mathrm{E}^{\mathfrak{F}}_\lambda \mu^{\{y\}} - \mu^{\{y\}}\right| + \mathrm{E}_\lambda(\mu^{B^{\mathsf{c}}}) < \epsilon.
\end{aligned}$$

□

Proof of Theorem 2.7: By Propositions 5.1 and 5.2, we immediately have $\lambda\mu \ll \lambda\nu$. For any $C \in \mathfrak{X} \times \mathfrak{Y}$,

$$\begin{aligned}\int \mu_x(C_x)\lambda(\mathrm{d}x) &= \int \overline{\mu}_x(C)\lambda(\mathrm{d}x)\\ &= (\lambda\mu)(C)\\ &= \int_C \frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}\mathrm{d}(\lambda\nu)\\ &= \int \lambda(\mathrm{d}x)\int_C \frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}\mathrm{d}\overline{\nu}_x\\ &= \int \lambda(\mathrm{d}x)\int_{C_x} \frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}(x,y)\nu_x(\mathrm{d}y),\end{aligned}$$

where the last equality follows from Proposition 5.3. Taking $C = A \times B$ for any $A \in \mathfrak{X}$ and $B \in \mathfrak{Y}$, we thus have

$$\int_A \mu_x(B)\lambda(\mathrm{d}x) = \int_A \lambda(\mathrm{d}x)\int_B \frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}(x,y)\nu_x(\mathrm{d}y),$$

so that

$$\mu_x(B) = \int_B \frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}(x,y)\nu_x(\mathrm{d}y),$$

for λ-almost every x in $\mathcal{X}$, and for every such x,

$$\frac{\mathrm{d}\mu_x}{\mathrm{d}\nu_x}(y) = \frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}(x,y)$$

for ν_x-almost every y in $\mathcal{Y}$. If $\mathrm{d}\mu_x/\mathrm{d}\nu_x$ has a $(\mathfrak{X}\times\mathfrak{Y})$-measurable version f, then

$$f = \frac{\mathrm{d}(\lambda\mu)}{\mathrm{d}(\lambda\nu)}$$

for all (x,y) except a $\lambda\nu$-negligible set of points (Proposition 5.4). □

Proof of Corollary 2.8:

$$\begin{aligned}&\mathrm{d}(\lambda\mu,\lambda\nu)\\ &= \frac{1}{2}\int\left|\frac{\mathrm{d}\lambda\mu}{\mathrm{d}(\lambda\times(\mu+\nu))} - \frac{\mathrm{d}\lambda\nu}{\mathrm{d}(\lambda\times(\mu+\nu))}\right|\mathrm{d}(\lambda\times(\mu+\nu))\\ &= \frac{1}{2}\int\lambda(\mathrm{d}x)\int\left|\frac{\mathrm{d}\lambda\mu}{\mathrm{d}(\lambda\times(\mu+\nu))}(x,y)\right.\\ &\quad\left. - \frac{\mathrm{d}\lambda\nu}{\mathrm{d}(\lambda\times(\mu+\nu))}(x,y)\right|(\overline{\mu+\nu})_x(\mathrm{d}(x,y))\\ &= \frac{1}{2}\int\lambda(\mathrm{d}x)\int\left|\frac{\mathrm{d}\mu_x}{\mathrm{d}(\mu+\nu)_x} - \frac{\mathrm{d}\nu_x}{\mathrm{d}(\mu+\nu)_x}\right|\mathrm{d}(\mu+\nu)_x \qquad (6)\\ &= \int \mathrm{d}(\mu_x,\nu_x)\lambda(\mathrm{d}x) = \mathrm{E}_\lambda\,\mathrm{d}(\mu_x,\nu_x),\end{aligned}$$

where (6) follow from Theorem 2.7 and Proposition 5.3. □

Proof of Proposition 2.11: It is obvious by observing that $[f:g] = \pi_{\mathbb{P}}(f(x),g(x))$ with f, g, and $\pi_{\mathbb{P}}$ all measurable. □

Proof of Proposition 2.12: By definition, for every $x \in \mathcal{X}$, there is a number $t(x) \neq 0$ such that $f_1(x) = t(x)f_2(x)$ and $g_1(x) = t(x)g_2(x)$. Then it suffices to show that t is measurable. Let $A = \{x : f_2(x) \neq 0\}$. It is clear that

$$t = \frac{f_1}{f_2 1_A + 1_{A^{\mathsf{c}}}}1_A + \frac{g_1}{g_2 1_{A^{\mathsf{c}}} + 1_A}1_{A^{\mathsf{c}}},$$

which is measurable. □

Proof of Proposition 2.13: Since $[f:g]$ is admissible μ-almost everywhere, we have $f \geq 0$, $g \geq 0$, and $(f,g) \neq (0,0)$ μ-almost everywhere, so that $\mu^{\mathfrak{F}}f \geq 0$, $\mu^{\mathfrak{F}}g \geq 0$, and $(\mu^{\mathfrak{F}}f, \mu^{\mathfrak{F}}g) \neq (0,0)$ μ-almost everywhere, and therefore $\mu^{\mathfrak{F}}[f:g]$ is admissible μ-almost everywhere. □

Proof of Theorem 2.15: First note that

$$[\mathrm{d}\mu|_{\mathfrak{F}} : \mathrm{d}\nu|_{\mathfrak{F}}] = (\mu+\nu)^{\mathfrak{F}}[\mathrm{d}\mu : \mathrm{d}\nu],$$

which is admissible $(\mu+\nu)$-almost everywhere by Proposition 2.13. From Corollary 2.5, it follows that there is a finite sub σ-algebra $\mathfrak{C}$ of $\mathfrak{X}$ and a finite sub σ-algebra $\mathfrak{D}$ of $\mathfrak{Y}$ such that

$$\begin{aligned}&\xi\left\{\left|\frac{\mathrm{d}\mu|_{\mathfrak{F}}}{\mathrm{d}(\mu+\nu)|_{\mathfrak{F}}} - \frac{\mathrm{d}\mu}{\mathrm{d}(\mu+\nu)}\right| \geq \epsilon\right\}\\ &\quad= \xi\left\{\left|(\mu+\nu)^{\mathfrak{F}}\frac{\mathrm{d}\mu}{\mathrm{d}(\mu+\nu)} - \frac{\mathrm{d}\mu}{\mathrm{d}(\mu+\nu)}\right| \geq \epsilon\right\} < \epsilon\end{aligned}$$

for every σ-algebra $\mathfrak{F}$ satisfying $\mathfrak{C}\times\mathfrak{D} \subseteq \mathfrak{F} \subseteq \mathfrak{X}\times\mathfrak{Y}$. Therefore

$$\begin{aligned}&\xi\{d_{\mathbb{P}}([\mathrm{d}\mu|_{\mathfrak{F}} : \mathrm{d}\nu|_{\mathfrak{F}}], [\mathrm{d}\mu : \mathrm{d}\nu]) \geq \epsilon\}\\ &\quad= \xi\left\{\left|\frac{\mathrm{d}\mu|_{\mathfrak{F}}}{\mathrm{d}(\mu+\nu)|_{\mathfrak{F}}} - \frac{\mathrm{d}\mu}{\mathrm{d}(\mu+\nu)}\right| \geq \epsilon\right\} < \epsilon.\end{aligned}$$

□

V. Facts Used By Section IV

Proposition 5.1: Let μ and ν be two kernels from the measure space $(\mathcal{X},\mathfrak{X},\lambda)$ to $(\mathcal{Y},\mathfrak{Y})$ such that $\mu_x \ll \nu_x$ for λ-almost every x in $\mathcal{X}$. Then $\lambda(\mu) \ll \lambda(\nu)$.

Proof: Let $B \in \mathfrak{Y}$. If $\lambda(\nu^B) = 0$, then $\nu(x,B) = 0$ for λ-almost every x in $\mathcal{X}$, so that $\mu(x,B) = 0$ for λ-almost every x in $\mathcal{X}$, hence $\lambda(\mu^B) = 0$, and therefore $\lambda(\mu) \ll \lambda(\nu)$. □

Proposition 5.2: Let μ and ν be two kernels from $(\mathcal{X},\mathfrak{X})$ to $(\mathcal{Y},\mathfrak{Y})$. For every $x \in \mathcal{X}$, $\mu_x \ll \nu_x$ iff $\overline{\mu}_x \ll \overline{\nu}_x$.

Proof: Let $B \in \mathfrak{Y}$ and $C \in \mathfrak{X}\times\mathfrak{Y}$.

($\Rightarrow$) If $\overline{\nu}(x,C) = 0$, then $\nu(x,C_x) = 0$, so that $\mu(x,C_x) = 0$, hence $\overline{\mu}(x,C) = 0$, and therefore $\overline{\mu}_x \ll \overline{\nu}_x$.

($\Leftarrow$) If $\nu(x,B) = 0$, then $\overline{\nu}(x,\mathcal{X}\times B) = 0$, so that $\overline{\mu}(x,\mathcal{X}\times B) = 0$, hence $\mu(x,B) = 0$, and therefore $\mu_x \ll \nu_x$. □

Proposition 5.3: Let f be a real-valued measurable function on $(\mathcal{X}\times\mathcal{Y},\mathfrak{X}\times\mathfrak{Y})$. Then $\overline{\mu}_x f = \mu_x f_x$.

Proof: Let ι_x be the map of $\mathcal{Y}$ into $\mathcal{X}\times\mathcal{Y}$ given by $y \mapsto (x,y)$, which is clearly measurable. Then

$$\overline{\mu}_x(f) = (\mu_x \circ \iota_x^{-1})f = \mu_x(f\circ\iota_x) = \mu_x f_x,$$

where the second equality follows from [3, Lemma 2.6.8]. □

Proposition 5.4: Let μ be a kernel from the measure space $(\mathcal{X},\mathfrak{X},\lambda)$ to $(\mathcal{Y},\mathfrak{Y})$. Let $A \in \mathfrak{X}$ and $C \in \mathfrak{X}\times\mathfrak{Y}$. If $\lambda(A) = 0$ and $\mu(x,C_x) = 0$ for all $x \notin A$, then $(\lambda\mu)(C) = 0$.

Proof:

$$\begin{aligned}(\lambda\mu)(C) &= (\lambda\mu)(C\cap(A^{\mathsf{c}}\times\mathcal{Y})) + (\lambda\mu)(C\cap(A\times\mathcal{Y}))\\ &\leq \int \overline{\mu}(x, C\cap(A^{\mathsf{c}}\times\mathcal{Y}))\lambda(\mathrm{d}x) + (\lambda\mu)(A\times\mathcal{Y})\\ &= \int_{A^{\mathsf{c}}} \mu(x,C_x)\lambda(\mathrm{d}x) + \lambda(A) = 0.\end{aligned}$$

□

VI. Conclusion

In this paper, we develop a general approach for deriving one-shot bounds for information-theoretic problems on general alphabets. This approach provides a mechanical way for solving problems on general alphabets based on their solutions in the finite-alphabet case, and hence it helps us better understand information theory in a unified way beyond countable alphabets. This is still an ongoing research. Applying this approach to other problems of information theory will be our future work.

Acknowledgment

This work was supported in part by the National Natural Science Foundation of China under Grant 61571398, Grant 61571006, and Grant 61371094, in part by the National Key Basic Research Program of China under Grant 2012CB316104, in part by the Zhejiang Provincial Natural Science Foundation under Grant LR12F01002, and in part by the open project of Zhejiang Provincial Key Laboratory of Information Processing, Communication and Networking.

References

[1] T. Han and S. Verdu, "Approximation theory of output statistics," *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 752–772, May 1993.

[2] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Hoboken, N.J: Wiley-Interscience, 2006.

[3] D. L. Cohn, *Measure Theory*, 2nd ed. Boston: Birkhäuser, 2013.

[4] R. Durrett, *Probability: Theory and Examples*, 4th ed., ser. Cambridge series in statistical and probabilistic mathematics. Cambridge ; New York: Cambridge University Press, 2010.

[5] O. Kallenberg, *Foundations of Modern Probability*, ser. Probability and its applications. New York: Springer, 1997.

[6] S. Yang, T. Honold, and Z. Zhang, "Beyond countable alphabets: An extension of the information-spectrum approach," extended version. [Online]. Available: http://arxiv.org/abs/1602.00376

[7] P. R. Halmos and L. J. Savage, "Application of the Radon-Nikodym theorem to the theory of sufficient statistics," *The Annals of Mathematical Statistics*, vol. 20, no. 2, pp. 225–241, Jun. 1949.

[8] J. R. Munkres, *Topology*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, Inc, 2000.

[9] M. H. Yassaee, M. R. Aref, and A. Gohari, "Achievability proof via output statistics of random binning," *IEEE Trans. Inf. Theory*, vol. 60, no. 11, pp. 6760–6786, Nov. 2014.

[10] S. Yang, "Separate random number generation from correlated sources," *IEEE Trans. Inf. Theory*, submitted for publication. [Online]. Available: http://arxiv.org/abs/1409.1662

(Version 0.5.1.20be8d)